arXiv:1611.01553v3 [cs.LO] 25 Nov 2016

QBF Solving by Counterexample-guided Expansion*

Roderick Bloem, Nicolas Braud-Santoni, and Vedad Hadzic

IAIK, Graz University of Technology, Austria

Abstract. We introduce a novel generalization of Counterexample-Guided Inductive Synthesis (CEGIS) and instantiate it to yield a novel, competitive algorithm for solving Quantified Boolean Formulas (QBF).
Current QBF solvers based on counterexample-guided expansion use a recursive approach which scales poorly with the number of quantifier alternations. Our generalization of CEGIS removes the need for this recursive approach, and we instantiate it to yield a simple and efficient algorithm for QBF solving.
Lastly, this research is supported by a competitive, though straightforward, implementation of the algorithm, making it possible to study the practical impact of our algorithm design decisions.

1 Introduction

Quantified Boolean Formulas (QBFs) are propositional logic formulas with existential and universal quantification. We consider QBF written in the prenex normal form $\exists x_0.\, \forall y_0.\, \ldots \forall y_n.\, \exists x_{n+1}.Q(x_0, \ldots, x_{n+1}, y_0, \ldots, y_n)$, with x_i and y_i representing disjoint sets of Boolean variables, and Q a propositional formula; any QBF formula can be easily and efficiently converted to prenex normal form.

Just as deciding satisfiability of a propositional-logic formula is the prototypical NP-complete problem, QBF satisfiability is the prototypical PSPACE-complete problem. Improvements in QBF solving will thus have a strong impact in many application domains with problems that can be translated to QBF, including planning [30], model checking [21], circuit design [25] and synthesis [7,32], in particular synthesis of distributed systems [15,16]. Research into engineering QBF solvers has seen considerable progress since initial work in the mid-nineties [10,22], with most solvers either implementing a quantified version of CDCL (e.g., [11,24,33]) or an expansion-based approach (e.g. [4]).

Since the first complete solver in 1998, based on DPLL [10], an incredible amount of work has gone into QBF: an analogue of propositional CDCL was discovered in 2002 [17, 23], the very same year as expansion-based solving [2] was discovered. Expansion-based solving has known tremendous success, and is the basis for counterexample-guided methods such as RAReQS [20], a state-of-the-art expansion-based solver. In 2016, a completely new approach to QBF solving, based on Skolem

* This work was supported by the Austrian Science Fund (FWF) through the Logic in Computer Science doctoral college (W1255-N23) and project RiSE (S11406-N23).

function computation, was introduced [28]. Still, expansion-based methods remain prevalent.

In this paper, we follow an expansion-based approach to solving QBFs that is inspired by counterexample-guided inductive synthesis (CEGIS) [1, 32]. CEGIS attempts to learn an input-output relation using a *learning algorithm* interacting with a *verification oracle*. It keeps a database of examples of input-output pairs. The learning algorithm proposes a candidate relation that is consistent with the examples in the database and sends it to the oracle (or gives up if no such relation can be found). The oracle adds counterexample to the candidate to the database (or declares that the proposed relation is correct).

We describe a novel relaxation of CEGIS that relies on an *approximate verifier*, *i.e.* a verifier which may produce spurious counterexamples: this makes the verifier simpler and faster. In exchange, the learner must provide values which prove that the candidate solution fulfils all the counterexamples produced so far.

The advantage of our approach is two-fold. First, it is extremely simple: the description of the algorithm is extremely small, and so is its proof of correctness. Second, it is efficient: our prototype implementation does not rely on any formula optimisation or implementation technique not described in this paper, yet it is competitive with state-of-the-art QBF solvers.

Our approach is similar to the one implemented in the very successful solver RAReQS [20], which instantiates the outer existential quantifiers using a counterexample-guided approach. When solving a formula of the form $\exists x_0 \forall y_0 \dots \forall y_n \exists x_{n+1}.Q$, RAReQS obtains the counterexample using a recursive call for $\exists y_0 \forall x_1 \dots \exists y_n \forall x_{n+1}.\neg Q$. To show that our approach is not simply a reformulation of RAReQS, we introduce in Section 4.3 a class of simple formulas on which RAReQS needs an exponential number of SAT queries while our approach needs only a small, constant number.

We have implemented our approach in a tool called *alejtehad*. We show that *alejtehad* is competitive with RAReQS, one of the best current QBF solvers, and that it is significantly outperforms RAReQS in some classes of benchmarks, especially those with high quantifier alternation depth.

2 Counterexample-Guided Inductive Synthesis (CEGIS)

2.1 Concepts

Counterexample-guided inductive synthesis (CEGIS) is a generic framework, initially devised in the context of syntax-guided synthesis [1], involving iterations between two components:

- The verifier provides, given a candidate solution, a counterexample that disproves it, or (correctly) proves the solution to be valid.
- The learner generates candidate solutions that are consistent with all counterexamples provided so far, or provides a (correct) unrealisability answer.

Although the worst-case complexity guarantees usually devolve to the size of the set of counterexamples or of candidate solution, this approach has been wildly successful in practice [20, 29, 32].

In a QBF setting, the standard approach is to consider that the solutions to be learned are values for the first quantifier; they must make the formula true, if the quantifier is $\exists$, or falsify it in the case of the $\forall$ quantifier.

However, one issue surfaces immediately when dealing with problems involving multiple quantifier alternations: the problem solved by the verifier is only marginally easier than the original one in practice.

For example, $\exists x.\, \forall y.\, \exists z.\, Q(x, y, z)$ is a type of QBF problem that occurs in synthesis of resilient systems, program diagnosis and repair, finite-state games, ... There, the verifier's problem is almost as hard as the original problem:

$$\exists y.\, \forall z.\, \neg\, Q(x, y, z)$$

One of the earliest applications of CEGIS[1] to QBF solving, RAReQS [20], introduced a recursive approach where the whole solver is used as a verifier, each nested solver dealing with one block of quantifiers.

There has been many refinements of this technique since RAReQS, from RAReQS' recent reimplementation from scratch for 2-QBF [3] to the clause abstraction performed by CAQE [29]. However, all of those share RAReQS' recursive approach to deal with quantifier alternation. This is an issue for efficiency, for multiple reasons:

- Performance degrades quickly as the quantifier depth increases, even on trivial instances, as shown in Section 4.3 and Section 5.
- Incremental solving, where SAT problems are constructed lazily by adding new conjuncts and solved without starting from scratch, becomes exceedingly difficult when sub-formulas change based on decisions made at arbitrary depths in the nested solvers.
- The solver does not share information deduced at different quantifier depths, resulting in duplicate work.
- The solver is forced, at any given point, to make decisions on variables belonging to a specific quantifier depth; we exploit this in Theorem 3 to highlight algorithmic differences between our approach and RAReQS.

Alternatively, the candidate solutions that are being discovered could be *Skolem functions* describing the values taken by the existentially-quantified variables. In that case, the natural form for counterexamples are *traces*: assignment of values to the universally-quantified variables that falsify the formula (under a given choice of Skolem functions).

However, this introduces a major difficulty: synthesis of Boolean (Skolem) functions is a hard problem on its own, that in general cannot be efficiently reduced to SAT (or a similarly-efficient decision procedure). Moreover, this would replace a situation where the learning is easy and the verification hard, with a problem where the learning is hard and the verification easy.

[1] Many of those are described in terms of Counterexample-Guided Abstraction Refinement (CEGAR), despite CEGIS being a closer fit to the algorithms.

2.2 Contribution

We take a slightly different, but related, approach: a candidate is a set E of *witnesses*, *i.e.* assignments to the existentially quantified variables; the learner's claim is that for any counterexample (*i.e.* any assignment of the universally quantified variables), the formula can be made true by instantiating the existentially-quantified variables to some element of E. Conversely, the database of counterexamples is a set A of assignments to all universally-quantified variables.

This means that the algorithm is symmetric between the universal and the existential variables: in order to find a new set of counterexamples (resp. candidate solution), we expand the formula using the candidate (resp. set of counterexamples) and call a SAT solver to find an assignment to the universal (resp. existential) variables. If the formula can be decided using a small set of candidates or a small set of counterexamples, this approach is fast.

This can be thought of as a relaxation of CEGIS in the sense that we allow an *approximate verifier*, *i.e.* a verifier which might produce spurious counterexamples. To appropriately deal with this, we must require that the learner not only produces a solution (an assignment to the outermost quantifier), but also auxiliary values (the witnesses) that prove that the proposed solution is consistent with the known counterexamples.

More explicitly, in the previous example ($\exists x.\, \forall y.\, \exists z.\, Q(x, y, z)$) the learner's problem becomes, at the nth iteration

$$\exists x, z_1 \dots z_n.\ \wedge_{i=1}^{n} Q(x, y_i, z_i)$$

while the verifier's problem becomes

$$\exists y_{n+1}.\ \wedge_{i=1}^{n} \neg Q(x, y_{n+1}, z_i)$$

While this enables us to iterate more rapidly between two simpler components, this comes at a cost: the solver may iterate over a correct value for x several times before providing enough witnesses for the verifier to determine that the solution is correct.

We show in Section 5 that, in practice, we outperform "traditional" CEGIS approaches on certain classes of problems, while being slower on others. This shows that this is a very promising approach: our algorithm is extremely simple and so are the associated proofs; yet the prototype implementation, written in less than 2 months, can compete with state-of-the-art solvers and outperforms them in some classes of problems, without using any optimisation not described in this paper.

3 A novel algorithm for QBF solving

We give a particular instantiation of this approach for QBF solving. We indifferently call our algorithm and its prototype implementation *alejtehad*, after the eponymous concept in Islamic philosophy: it refers both to the argumentation

process through which case law is produced, and to the state of knowing that one knows nothing.

This hints at the minimalistic nature of the algorithm: *alejtehad* is simple, in the sense that it is concise to describe and has few prerequisites, both in the theory and in its implementation. It relies on an iterative process, in which arguments and counterarguments are produced until a solution is found.

Beyond the novelty and simplicity of its solving algorithm, *alejtehad* is unusual in that it efficient as described, without relying on complex formula optimisations or implementation strategies not described in Section 3.3.

3.1 Notation

In the entirety of this paper, Φ will denote a QBF formula and Q the *matrix* (*i.e.* the propositional body) of a QBF formula. We consider, without loss of generality, prenex formulas where each quantifier ranges over a single bitvector:

$$\Phi := \exists x_0. \forall y_0. \ldots \forall y_n. \exists x_{n+1}. Q\left(x_0, \ldots x_{n+1}, y_0 \ldots y_n\right)$$

We denote a single bitvector by z_i, a tuple of bitvectors by Z, and their assigned values (in a given model) by $\mathbf{z}_i$ and $\mathbf{Z}$. Under this notation, we can now precisely define witnesses and counterexamples:

- Witnesses are tuples $\mathbf{X} := (\mathbf{x}_0, \ldots, \mathbf{x}_{n+1})$ of concrete values for the existentially-quantified variables. $\mathbf{X}$ is said to be a witness for a counterexample $\mathbf{Y}$ if $Q\left(\mathbf{X}, \mathbf{Y}\right)$ holds.
- Counterexamples are tuples $\mathbf{Y} := (\mathbf{y}_0, \ldots, \mathbf{y}_n)$ of concrete values for the universally-quantified variables. $\mathbf{Y}$ is a counterexample of a witness $\mathbf{X}$ if $\neg\, Q(\mathbf{X}, \mathbf{Y})$ holds.
- Witness variables are tuples $X^{\mathbf{Y}} := \left(x_0, x_1^{\mathbf{y}_0}, x_2^{\mathbf{y}_0\mathbf{y}_1}, \ldots x_{n+1}^{\mathbf{y}_0 \cdots \mathbf{y}_n}\right)$ associated to a given counterexample $\mathbf{Y}$. Given a model, $\mathbf{X}^{\mathbf{Y}}$ denotes its concrete value (in a given model), *i.e.* a *witness*.
- Counterexample variables are likewise defined, denoted $Y^{\mathbf{X}}$, and their concrete values, denoted $\mathbf{Y}^{\mathbf{X}}$, are counterexamples.

For example, for the QBF $\Phi := \forall y. \exists x. Q(x, y)$ with $Q(x, y) := x = y + 1$ (interpreting x and y as n-bit integers), the counterexample 0 has a corresponding witness variable x^0. By solving[2] $Q(x^0, 0)$ (*i.e.* $x^0 = 1$), a matching witness $\mathbf{x}^0 = 1$ is discovered, and has a corresponding counterexample variable y^1. Solving the SAT instance $\neg Q(\mathbf{x}^0, y^{\mathbf{x}^0})$ then yields a new counterexample $\mathbf{y}^1$.

If a prenex QBF formula as defined above is true, it has an associated *Skolem model* [31]: functions $f_{x_{i+1}}(y_0 \ldots y_i)$ such that, for any counterexample $\mathbf{Y}$, $(f_{x_0}(), f_{x_1}(\mathbf{y}_0), \ldots)$ is a matching witness. Conversely, false formulas have *Herbrand models* [18]: functions $f_{y_i}(x_0 \ldots x_i)$ such that for any witness $\mathbf{X}$, $(f_{y_0}(\mathbf{x}_0), f_{y_1}(\mathbf{x}_0, \mathbf{x}_1) \ldots)$ is a matching counterexample.

Thinking of the QBF formula as a 2-players game, a Skolem (resp. Herbrand) model is a winning strategy for the existential (resp. universal) player. We use the

[2] This is a propositional SAT problem.

notions of Skolem and Herbrand models later, in Section 4, while proving that *alejtehad* is sound.

3.2 The *alejtehad* algorithm

As described previously, our algorithm is composed of a learner and an approximate verifier. It maintains a database of counterexamples (denoted A) and a database of witnesses (denoted E).

At each iteration, the learner produces witnesses for all counterexamples observed so far, at which time the verifier responds with counterexamples for all witnesses; both consist of a single query to Boolean SAT. The solver keeps iterating until the learner or the verifier returns UNSAT, at which point the solver returns either a value for x_0, if there are no counterexamples, or false if there are no satisfying witnesses.

A more precise description is provided in Algorithm 1, while Section 3.3 exposes the differences between this description and an efficient implementation.

Result: false or a satisfying value $\mathbf{x}_0$
Initially, $E := \emptyset$ and $A := \{\mathbf{Y}_0\}$ contains a single, arbitrary counterexample;
while *true* **do**
  `// Learner`
  if $\varphi := \bigwedge_{\mathbf{Y} \in A} Q\left(X^{\mathbf{Y}}, \mathbf{Y}\right)$ *is satisfiable* **then**
    Add each witness $\mathbf{X}^{\mathbf{Y}}$ to E;
  else
    return false;
  end
  `// Verifier`
  if $\psi := \bigwedge_{\mathbf{X} \in E} \neg Q\left(\mathbf{X}, Y^{\mathbf{X}}\right)$ *is satisfiable* **then**
    Add each counterexample $\mathbf{Y}^{\mathbf{X}}$ to A;
  else
    return $\mathbf{x}_0$;
  end
end

Fig. 1. *ijtihad* algorithm

3.3 Implementation concerns

Obviously, Algorithm 1 only gives a high-level description of the solving algorithm, which isn't quite the description of a fast implementation. Here are notable aspects of our implementation that makes it perform well:

- At each iteration, new propositional conjuncts are introduced in φ and ψ: we systematically use incremental SAT solving.

– Existing clauses aren't introduced in duplicate in φ or ψ.
– Likewise, when converting $\neg Q$ to CNF, naively using Tseitin's results in a large (though linear) blowup: instead, we maintain a cache that associates to each clause a Tseitin variable; when a clause is encountered was already processed, no new clauses are added to ψ and the existing Tseitin variable is reused.

While RAReQS, against which we compare our implementation in Section 5 uses similar techniques in its implementation, it also includes optimisations performed on the QBF formula, both before and during solving, that we do not implement: unit propagation, pure literals elimination, universal reduction and blocking clauses. Indeed, *alejtehad* does not consider explicitly QBF subproblems, so it is unclear whether existing QBF simplification techniques can be applied thorough the execution of the algorithm.

4 Formal arguments

As mentioned earlier, Algorithm 1 is very simple and so are its termination and soundness arguments. We provide no complexity guarantees beyond the immediate upper-bound that follows from the termination argument (Theorem 1): while counterexample-guided methods perform well in practice, they do so at the cost of worst-case bounds.

However, we exhibit certain classes of problems that *alejtehad* can deal within 3 SAT queries, whereas existing CEGIS-based methods, represented by RAReQS, require exponentially many.

4.1 Termination

Lemma 1. *At every iteration of Algorithm 1, A increases strictly (with regards to set inclusion).*

Proof. The only operation affecting A is adding to it, so A increases.

Hence, it is sufficient to prove that at least one new counterexample is produced at each iteration. For a given A, if $\Phi := \bigwedge_{\mathbf{Y} \in A} Q\left(X^{\mathbf{Y}}, \mathbf{Y}\right)$ is satisfiable, $Q(X^{\mathbf{Y}}, \mathbf{Y})$ holds in the resulting model, for each counterexample $\mathbf{Y} \in A$.

Hence, $\psi := \wedge_{\mathbf{X} \in E} \neg Q\left(\mathbf{X}, Y^{\mathbf{X}}\right)$ cannot be satisfied using only existing counterexamples (concrete values that are already in A): either ψ is unsatisfiable (and the algorithm terminates) or the produced model contains new values that are added to A, which thus grows (strictly) larger.

Q.e.d. □

Theorem 1. *Algorithm 1 terminates after at most $2^{1+\sum_{i=0}^{n}|y_i|}$ SAT queries.*

Proof. The set of all possible counter-examples, a superset of A, is finite and of size $2^{\sum_{i=0}^{n}|y_i|}$.

From Lemma 1, at least one new counterexample is produced at each iteration: there are at most $2^{\sum_{i=0}^{n}|y_i|}$ iterations.

Lastly, two SAT queries are performed at each iteration, except for the last which may perform one or two queries.

Q.e.d. □

4.2 Soundness

Lemma 2. *If the QBF formula Φ is true, Algorithm 1 returns some $\mathbf{x}_0$ such that Φ admits a Skolem model of the formula where the value of x_0 is $\mathbf{x_0}$.*

Proof. If the original formula is true, there is a Skolem model [31] for it:

$\mathbf{f}_{x_0}(), \mathbf{f}_{x_1}(\cdot), \ldots \mathbf{f}_{x_{n+1}}(\cdot)$ such that $\forall y_0. \ldots y_n. Q\left(\mathbf{f}_{x_0}(), y_0, \mathbf{f}_{x_1}(y_0), \ldots y_{n-1}, \mathbf{f}_{x_n}(y_0, \ldots y_n)\right)$

Then, $\varphi := \bigwedge_{\mathbf{Y} \in A} Q\left(X^{\mathbf{Y}}, \mathbf{Y}\right)$ is always satisfiable: $x_{i+1}^{\mathbf{y}_0 \ldots \mathbf{y}_i} := \mathbf{f}_{x_{i+1}}(\mathbf{y}_0, \ldots \mathbf{y}_i)$ (for each $(\mathbf{y}_0, \ldots \mathbf{y}_n) \in A$) is a satisfying model. Hence, Algorithm 1 cannot return false, and it terminates (Theorem 1), so it must returns some $\mathbf{x}_0$.

Let us consider the following formula:

$$\forall y_0. \exists x_1 \forall y_1 \ldots . \exists x_{n+1}. Q(\mathbf{f}_{x_0}, x_1, \ldots x_{n+1}, Y)$$

If this formula is true, then adding $x_0 := \mathbf{f}_{x_0}$ to one of its Skolem models yields a Skolem model of the original QBF with the desired property.

If it is false, then a model for $\psi := \bigwedge_{\mathbf{X} \in E} \neg Q\left(\mathbf{X}, Y^{\mathbf{X}}\right)$ can be constructed from its Herbrand model, using the same construction as in Lemma 3; thus, ψ is satisfiable and the algorithm cannot have returned at that point.

Q.e.d. □

Lemma 3. *If the QBF formula is false, Algorithm 1 returns* false.

Proof. If the original QBF is false, it has a Herbrand model [18]: $\mathbf{f}_{y_0}(\cdot) \ldots \mathbf{f}_{y_n}(\cdot)$ such that, for each $x_1 \ldots x_{n+1}$, $\neg Q\left(x_0, \mathbf{f}_{y_0}(x_0), \ldots x_n, \mathbf{f}_{y_n}(x_0 \ldots x_n), x_{n+1}\right)$

Hence, ψ has a satisfying model, $y_i^{\mathbf{x}_0 \ldots \mathbf{x}_i} := \mathbf{f}_{y_i}(\mathbf{x}_0 \ldots \mathbf{x}_i)$ (for each $(\mathbf{x}_0, \ldots \mathbf{x}_{n+1}) \in A$), so Algorithm 1 cannot return some $\mathbf{x}_0$.

Since it must return (Theorem 1), it must then return false.

Q.e.d. □

The soundness result follows immediately:

Theorem 2. *The result returned by Algorithm 1 is consistent with the value of the input..*

4.3 Distinguishability from RAReQS

As mentioned earlier, counterexample-guided methods usually an have exponential worst-case complexity, which is disconnected from their performance in practice. This makes it difficult to study those methods through the lens of complexity theory.

However, we can give here a family of examples which highlights the algorithmic differences between our approach and "traditional" CEGIS-based solvers, such as RAReQS.

Note that we compare against the algorithmic description of RAReQS, without considering additional optimizations, especially in-processing, which are not described in the RAReQS paper.

Theorem 3. *There exists an infinite family of QBF terms such that ijtihad performs $\mathcal{O}(1)$ SAT queries solving them, whereas RAReQS' algorithm performs $\Omega(2^n)$ queries.*

We define those formulas to have the following structure, with Q_n^X satisfiable and Q_n^Y falsifiable, X being $(x_0, \dots x_{n+1})$ and Y being $(y_0, \dots y_n)$:

$$\Phi_n := \exists x_0. \forall y_0. \dots \forall y_n. \exists x_{n+1}.\ Q_n^X(X) \wedge Q_n^Y(Y)$$

Lemma 4. *ijtihad performs either 1 or 3 SAT queries to solve Φ_n, for any n.*

Proof.
1. $\mathbf{Y}_0$ is taken arbitrarily.
2. A first SAT call is performed to obtain a witness $\mathbf{X}_0$.
   - If $\mathbf{Y}_0$ falsifies Q_n^Y, the result is UNSAT, and ijtihad returns.
   - Otherwise, $\mathbf{X}_0$ is such that $Q_n^X(\mathbf{X}_1)$ holds.
3. Since $Q_n^X(\mathbf{X}_0)$ holds, the counterexample $\mathbf{Y}_1$ must falsify Q_n^Y.
4. The next SAT call is guaranteed to yield UNSAT, as no witness exists.

Q.e.d. □

For convenience, we define auxiliary families of formulas:

$$\begin{aligned}
\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i,\mathbf{y}_i)} &:= \exists x_{i+1}. \forall y_{i+1}. \exists x_{i+2} \dots \exists x_{n+1}. \\
&\qquad Q_n^X(\mathbf{x}_0, \dots \mathbf{x}_i, x_{i+1}, \dots x_{n+1}) \wedge Q_n^Y(\mathbf{y}_0, \dots \mathbf{y}_i, y_{i+1}, \dots y_n) \\
\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i)} &:= \exists y_i. \forall x_{i+1}. \exists y_{i+1} \dots \forall x_{n+1}. \\
&\qquad \neg\, Q_n^X(\mathbf{x}_0, \dots \mathbf{x}_i, x_{i+1}, \dots x_{n+1}) \vee \neg\, Q_n^Y(\mathbf{y}_0, \dots \mathbf{y}_{i-1}, y_i, \dots y_n)
\end{aligned}$$

Note that $\Phi_n = \Phi_n^{()}$ and $\neg\, \Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i,\mathbf{y}_i)} = \forall x_{i+1}. \Phi_n^{\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i,\mathbf{y}_i,x_{i+1}}$.

Lemma 5. *RAReQS' algorithm performs $\Omega\left(2^{n-i}\right)$ SAT queries when called on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i,\mathbf{y}_i)}$.*

Proof. The execution of RAReQS on $\Phi_n^{(\mathbf{X},\mathbf{Y})}$ results in a single SAT call, as there is only a single quantifier.

The execution of RAReQS on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i,\mathbf{y}_i)}$ goes as follows [20]:

1. An arbitrary value $\mathbf{x}_{i+1}$ is picked. Let us assume[3] that is leaves Q_n^X satisfiable.
2. RAReQS is recursively called on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\dots\mathbf{x}_i,\mathbf{y}_i,\mathbf{x}_{i+1})}$

[3] This is the best case for RAReQS, but it doesn't matter for the sake of our argument.

(a) An arbitrary value $\mathbf{y}_{i+1}$ is picked. Let us assume it leaves Q_n^Y not falsifiable.
(b) RAReQS is recursively called on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\ldots\mathbf{x}_i,\mathbf{y}_i,\mathbf{x}_{i+1},\mathbf{y}_{i+1})}$, yielding $\mathbf{x}_{i+1}$.
(c) RAReQS is recursively called on $\exists y_{i+1}.\neg\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\ldots\mathbf{x}_i,\mathbf{y}_i,\mathbf{x}_{i+1},y_{i+1})}$, yielding $\mathbf{y}_{i+1}$ such that Q_n^Y is refutable.
(d) RAReQS is called on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\ldots\mathbf{x}_i,\mathbf{y}_i,\mathbf{x}_{i+1},\mathbf{y}_{i+1})}$, yielding false.
(e) $\mathbf{y}_{i+1}$ is returned from the recursive RAReQS call.

3. The execution of RAReQS on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\ldots\mathbf{x}_i,\mathbf{y}_i)}$ resumes.

RAReQS performs, while solving $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\ldots\mathbf{x}_i,\mathbf{y}_i)}$, at least two recursive calls on $\Phi_n^{(\mathbf{x}_0,\mathbf{y}_0,\ldots\mathbf{x}_{i+1},\mathbf{y}_{i+1})}$ with different values of $\mathbf{y}_{i+1}$: it follows that $\Omega\left(2^{n-i}\right)$ SAT calls are performed.

Q.e.d. □

5 Experimental results

In this section we look at the performance of *alejtehad* against RAReQS, which won several recent QBF competitions. We also point out different families of problems where we do better, as well as those where we do very poorly. Furthermore, we point out the impact of various formula parameters on our solving time.

5.1 Experimental setting

The optimised *alejtehad* algorithm is implemented in C++ and uses MiniSat [12–14] as its backing SAT solver. The code is based on RAReQS', so as to minimise the differences (data structures, cache optimisations, ...) that could explain performance differences, aside from the implemented algorithm. It is available, under a free software licence (GPLv3), at our laboratory's source code repository[4].

The input QBF is given in the standard QDIMACS format. All experiments were run on a dedicated Nehalem machine with 192GB of RAM and 24 CPU threads, running at 2.8GHz each. Each execution was allocated a single virtual core and 10 minutes per problem, while the machine was otherwise unused.

RAReQS was run using its default configuration, whereas our solver does not have tweakable parameters.

5.2 Performance on Application Benchmarks

We ran both solvers on the 825 benchmarks from the main track of the QBFEval'16[5] competition. We also evaluated the impact of the preprocessor Bloqqer [6,19], which yields an overall performance increase for both solvers but is more helpful to RAReQS, as seen in Table 1.

We then analysed how the number of quantifiers affects the performance of the solvers, which can be seen in Figure 2. It becomes clear that *alejtehad* tends

[4] `https://extgit.iaik.tugraz.at/scos/ijtihad`
[5] Available on `http://www.qbflib.org/eval16.html` at the time of writing.

Solver	Solved instances	SAT	UNSAT	Solving time (s) Total	Average
RAReQS + b	627	306	321	16316.9	26.0
alejtehad + b	545	261	284	5153.9	9.5
RAReQS	331	125	206	8505.9	25.7
alejtehad	322	124	198	8739.3	27.1

Table 1. Number of solved test cases, time needed for solving them, and the average solving time of solved test cases for each of the solvers (not including timeouts).

to have an advantage over RAReQS for QBFs with more quantifier alternations. Using Bloqqer, however, generally decreases the number of quantifiers [6], which explains why preprocessing helps RAReQS more.

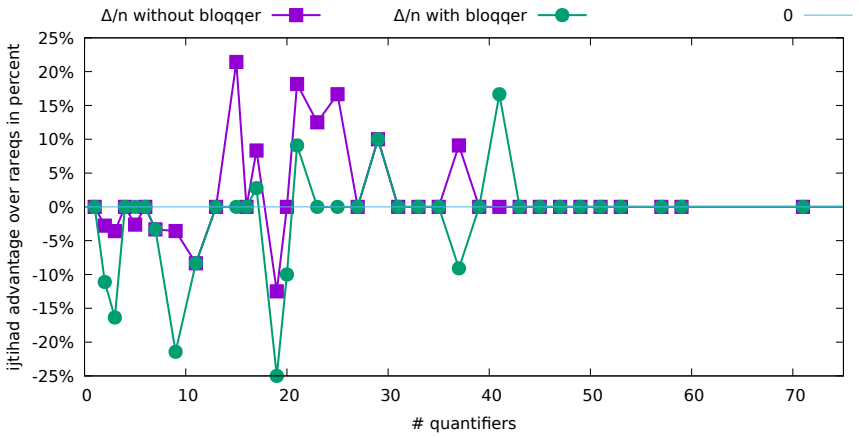

Fig. 2. Plot showing the difference in number of solved test cases by depth[6]

The biggest difference in performance can be seen on the benchmark families described by Pan and Vardi [27], and those described by Mneimneh and Sakallah [26]. As shown in Figure 3 and Table 2, *alejtehad* is faster on the Pan family, which has higher quantifier depth. On the other hand, RAReQS outperforms it on the Mneimneh-Sakallah family, which exhibits many variables and low quantification depth. In fact, one could argue that the Mneimneh-Sakallah QBF family is the main reason RAReQS' results seem better in Table 1, as *alejtehad* almost always runs into a timeout in this specific class, which is overrepresented, accounting for approximately 12% of the entire set of benchmarks.

[6] Only quantifier depths with more than one test case are shown in order to reduce noise.

Benchmarks	# Test Cases	Solver	# Solved	Solving time
Pan	181	RAReQS + b	171	879.681
		alejtehad + b	173	710.27
Mneimneh-Sakallah	100	RAReQS + b	73	9748
		alejtehad + b	12	323.19

Table 2. Number of solved test cases and time needed for solving them, for the QBF families Pan and Mneimneh-Sakallah

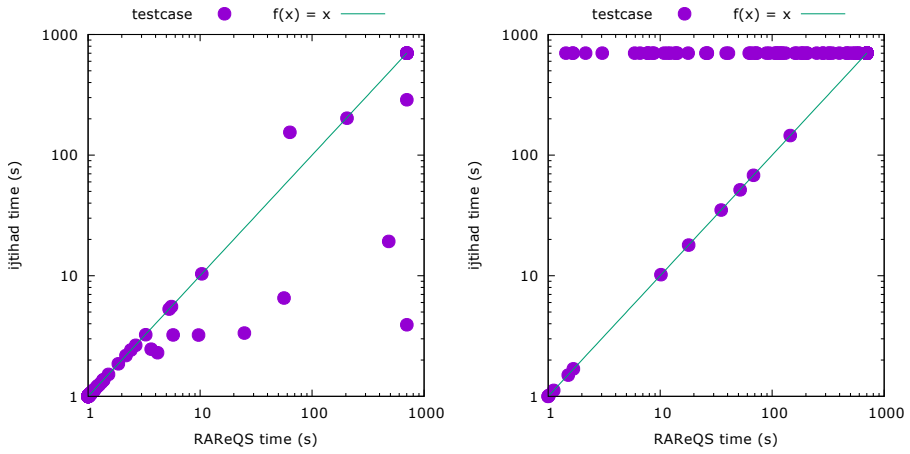

Fig. 3. Log-log time scatter plots of the Pan family on the left, and Mneimneh-Sakallah family on the right. *alejtehad*'s solving time is shown on the **y** axis, while RAReQS' is on the **x** axis.

5.3 Performance on Random QBFs with a fixed structure

Experimental evidence shown in [9] suggests that there is a crossover point in the size of random QBF problems, where the formula is equally likely to be SAT and UNSAT, and where verifying or falsifying it becomes as hard as possible. We use this to argue that *alejtehad* and RAReQS scale differently as quantifier depth increases.

In order to further support the claim that *alejtehad* outperforms RAReQS on QBFs with many quantifier alternations, we ran both solvers without external preprocessing[7] on hard random QBFs. All examined QBFs were generated by BlocksQBF [8] and had a fixed structure in which each clause of the quantified CNF had one variable from every quantified variable vector.

We produced benchmarks with increasing numbers of variables under each quantifier until both solvers ran into timeouts, and then compared the overall results. Figure 4 shows the performance of the solvers for the number of variables which produced the biggest difference in behaviour, for a given number of quantifiers.

As can be observed, on problems with more than 4 quantifiers, we consistently solve more problems than RAReQS, and faster. One important thing to note is that timeouts weight in the solving time, with no particular penalty beyond the timeout value; this explains that our speedup seems to taper off at quantification depths 12 to 14: *alejtehad* takes more time to solve instances, and so seems closer to RAReQS despite the fact that it doesn't solve any instances at all.

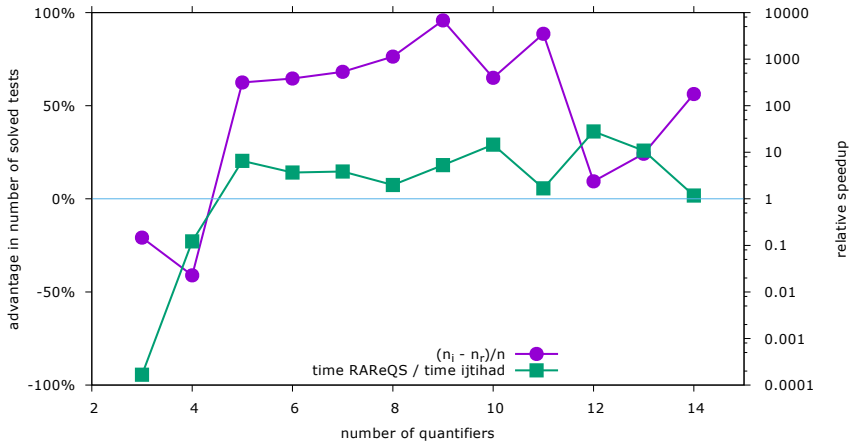

Fig. 4. Plot showing the difference in number of solved test cases by depth, as well as the relative time. In the graph, n_i and n_r denote the number of test cases solved by *alejtehad* and RAReQS respectively, whereas n denotes the total number of test cases (with the given quantifier depth).

[7] The number of quantifiers is important and Bloqqer would change it

Conclusion

To summarise our contribution, we provide a new perspective on Counterexample-Guided Inductive Synthesis (CEGIS), which yields a novel QBF-solving algorithm, which is both extremely simple to describe and prove correct, and which our experiments show to be efficient.

Those findings are supported by an implementation of this algorithm, *alejtehad*, using no optimisation or implementation techniques not described in the paper. This implementation is competitive with RAReQS, a state-of-the-art solver which won recent QBF competitions. In particular, Section 5 shows that our approach outperforms RAReQS in problems involving more than 4 nested quantifiers.

On the other hand, our approach shares a weakness with other expansion-based methods (including RAReQS): very simple formulas like $\forall x.\, \exists y.\, y = x$ take 2^n steps to enumerate the (x, y) pairs [28], with x and y being vectors of n Boolean variables. Alternative approaches may not suffer from this specific issue. Most notably, Rabe *et al.* have been investigating synthesising directly Skolem functions [28, 29], which can efficiently deal with this kind of QBF instances.

There are many venues of research this opens for future work:

- Are there commonplace QBF optimisations that are applicable here, and do they improve the performance of *alejtehad*?
- Preprocessors such as Bloqqer are designed for solvers for which deeper quantification is problematic. As exhibited in Section 5.3, current preprocessors do not benefit *alejtehad* as much as other solver: can Bloqqer be tuned or improved to better support our approach?
- Can our approach be combined with Skolem function synthesis [28] so as to handle problems where the Skolem or Herbrand models contain many concrete values?
- Outputting small certificates is a known challenge for counterexample-guided methods. More generally, can we extract small Skolem and Herbrand models?
- Can we generalise our implementation to NNF (Negation Normal Form) without losing efficiency? NNF is interesting, especially for QBF, as encoding problems into NNF does not require introducing new variables, and so does negating NNF formulas.
- How does search-based solving compare to our approach? Can we combine both approaches?
- The *alejtehad* algorithm is intrinsically sequential, due to its data dependencies. Can we find a variation on it that can be efficiently parallelised?

Acknowledgements

We would like to thank very warmly Armin Biere, Florian Lonsing and Martina Seidl, with whom we cooperate through the LogiCS doctoral college and the RiSE project, for their patient support and willingness to share their expertise in QBF

solving. Moreover, Martina Seidl and Uwe Egly deserve very special thanks for their expert and thorough review of the paper.

Our understanding of our relaxation of CEGIS also benefited from fruitful discussions at CAV'16 with Armando Solar-Lezama, the original discoverer of CEGIS.

References

1. Alur, R., Bodík, R., Dallal, E., Fisman, D., Garg, P., Juniwal, G., Kress-Gazit, H., Madhusudan, P., Martin, M. M. K., Raghothaman, M., Saha, S., Seshia, S. A., Singh, R., Solar-Lezama, A., Torlak, E., and Udupa, A. Syntax-guided synthesis. In *Dependable Software Systems Engineering*, M. Irlbeck, D. A. Peled, and A. Pretschner, Eds., vol. 40 of *NATO Science for Peace and Security Series, D: Information and Communication Security*. IOS Press, 2015, pp. 1–25.
2. Ayari, A., and Basin, D. *Qubos: Deciding Quantified Boolean Logic Using Propositional Satisfiability Solvers*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 187–201.
3. Balabanov, V., Jiang, J.-H. R., Scholl, C., Mishchenko, A., and Brayton, R. K. *2QBF: Challenges and Solutions*. Springer International Publishing, Cham, 2016, pp. 453–469.
4. Biere, A. Resolve and expand. In *Theory and Applications of Satisfiability Testing, 7th International Conference, SAT 2004, Vancouver, BC, Canada, May 10-13, 2004, Revised Selected Papers* (2004), H. H. Hoos and D. G. Mitchell, Eds., vol. 3542 of *Lecture Notes in Computer Science*, Springer, pp. 59–70.
5. Biere, A., Heule, M., van Maaren, H., and Walsh, T., Eds. *Handbook of Satisfiability*, vol. 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2009.
6. Biere, A., Lonsing, F., and Seidl, M. *Blocked Clause Elimination for QBF*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 101–115.
7. Bloem, R., Egly, U., Klampfl, P., Könighofer, R., and Lonsing, F. Sat-based methods for circuit synthesis. In *Formal Methods in Computer-Aided Design, FMCAD 2014, Lausanne, Switzerland, October 21-24, 2014* (2014), IEEE, pp. 31–34.
8. Brummayer, R., Lonsing, F., and Biere, A. Automated testing and debugging of sat and qbf solvers. In *In Proc. of SAT 2010, volume 6175 of LNCS* (2010), Springer, pp. 44–57.
9. Cadoli, M., Giovanardi, A., and Schaerf, M. Experimental analysis of the computational cost of evaluating quantified boolean formulae. In *In: AI*IA 97* (1997), Springer-Verlag, pp. 207–218.
10. Cadoli, M., Giovanardi, A., and Schaerf, M. An algorithm to evaluate quantified boolean formulae. *AAAI/IAAI 98* (1998), 262–267.
11. Cadoli, M., Schaerf, M., Giovanardi, A., and Giovanardi, M. An algorithm to evaluate quantified boolean formulae and its experimental evaluation. *J. Autom. Reasoning 28*, 2 (2002), 101–142.
12. Eén, N., and Biere, A. Effective preprocessing in sat through variable and clause elimination. In *Proceedings of the 8th International Conference on Theory and Applications of Satisfiability Testing* (Berlin, Heidelberg, 2005), SAT'05, Springer-Verlag, pp. 61–75.
13. Eén, N., Mishchenko, A., and Sörensson, N. Applying logic synthesis for speeding up SAT. In *Theory and Applications of Satisfiability Testing - SAT 2007, 10th International Conference, Lisbon, Portugal, May 28-31, 2007, Proceedings* (2007), J. Marques-Silva and K. A. Sakallah, Eds., vol. 4501 of *Lecture Notes in Computer Science*, Springer, pp. 272–286.
14. Eén, N., and Sörensson, N. *An Extensible SAT-solver*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 502–518.

15. Finkbeiner, B., and Tentrup, L. Detecting unrealizability of distributed fault-tolerant systems. *Logical Methods in Computer Science 11*, 3 (2015).
16. Gascón, A., and Tiwari, A. A synthesized algorithm for interactive consistency. In *NASA Formal Methods Symposium* (2014), Springer, pp. 270–284.
17. Giunchiglia, E., Narizzano, M., Tacchella, A., et al. Learning for quantified boolean logic satisfiability. In *AAAI/IAAI* (2002), pp. 649–654.
18. Herbrand, J. Investigations in proof theory: The properties of true propositions. in van heijenoort, j., editor. *From Frege to Gödel* (1967), 525–582.
19. Heule, M., Järvisalo, M., Lonsing, F., Seidl, M., Biere, A., et al. Clause elimination for sat and qsat. *JAIR 53* (2015), 127–168.
20. Janota, M., Klieber, W., Marques-Silva, J., and Clarke, E. M. Solving QBF with counterexample guided refinement. *Artif. Intell. 234* (2016), 1–25.
21. Jussila, T., and Biere, A. Compressing BMC encodings with QBF. *Electr. Notes Theor. Comput. Sci. 174*, 3 (2007), 45–56.
22. Kleine Büning, H., Karpinski, M., and Flögel, A. Resolution for quantified boolean formulas. *Inf. Comput. 117*, 1 (1995), 12–18.
23. Letz, R. Lemma and model caching in decision procedures for quantified boolean formulas. In *Proceedings of the International Conference on Automated Reasoning with Analytic Tableaux and Related Methods* (London, UK, UK, 2002), TABLEAUX '02, Springer-Verlag, pp. 160–175.
24. Lonsing, F., and Egly, U. Incrementally computing minimal unsatisfiable cores of qbfs via a clause group solver API. In *Theory and Applications of Satisfiability Testing - SAT 2015 - 18th International Conference, Austin, TX, USA, September 24-27, 2015, Proceedings* (2015), M. Heule and S. Weaver, Eds., vol. 9340 of *Lecture Notes in Computer Science*, Springer, pp. 191–198.
25. Mangassarian, H., Veneris, A., and Benedetti, M. Robust qbf encodings for sequential circuits with applications to verification, debug, and test. *IEEE Transactions on Computers 59*, 7 (July 2010), 981–994.
26. Mneimneh, M., and Sakallah, K. Computing vertex eccentricity in exponentially large graphs: Qbf formulation and solution. In *In Proc. 6th Intl. Conf. on Theory and Applications of Satisfiability Testing (SAT03), volume 2919 of LNCS* (2003), Springer, pp. 411–425.
27. Pan, G., and Vardi, M. Y. Optimizing a bdd-based modal solver. In *In Proceedings of the 19th International Conference on Automated Deduction* (2003), Springer, pp. 75–89.
28. Rabe, M. N., and Seshia, S. A. *Incremental Determinization*. Springer International Publishing, Cham, 2016, pp. 375–392.
29. Rabe, M. N., and Tentrup, L. Caqe: A certifying qbf solver. In *Proceedings of the 15th Conference on Formal Methods in Computer-Aided Design* (Austin, TX, 2015), FMCAD '15, FMCAD Inc, pp. 136–143.
30. Rintanen, J. Constructing conditional plans by a theorem-prover. *J. Artif. Intell. Res. (JAIR) 10* (1999), 323–352.
31. Skolem, T. Logico-combinatorial investigations in the satisfiability or provability of mathematical propositions: a simplified proof of a theorem by l. löwenheim and generalizations of the theorem. *From Frege to Gödel. A Source Book in Mathematical Logic 1931* (1879), 252–263.
32. Solar-Lezama, A. Program sketching. *STTT 15*, 5-6 (2013), 475–495.
33. Zhang, L., and Malik, S. *Towards a Symmetric Treatment of Satisfaction and Conflicts in Quantified Boolean Formula Evaluation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 200–215.